



Do code review measures explain the incidence of post-release defects?

Case study replications and bayesian networks

Andrey Krutauz¹ · Tapajit Dey² · Peter C. Rigby¹  · Audris Mockus²

Published online: 29 June 2020

© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Aim In contrast to studies of defects found during code review, we aim to clarify whether code review measures can explain the prevalence of post-release defects.

Method We replicate McIntosh *et al.*'s (Empirical Softw. Engg. **21**(5): 2146–2189, 2016) study that uses additive regression to model the relationship between defects and code reviews. To increase external validity, we apply the same methodology on a new software project. We discuss our findings with the first author of the original study, McIntosh. We then investigate how to reduce the impact of correlated predictors in the variable selection process and how to increase understanding of the inter-relationships among the predictors by employing Bayesian Network (BN) models.

Context As in the original study, we use the same measures authors obtained for Qt project in the original study. We mine data from version control and issue tracker of Google Chrome and operationalize measures that are close analogs to the large collection of code, process, and code review measures used in the replicated the study.

Results Both the data from the original study and the Chrome data showed high instability of the influence of code review measures on defects with the results being highly sensitive to variable selection procedure. Models without code review predictors had as good or better fit than those with review predictors. Replication, however, confirms with the bulk of prior work showing that prior defects, module size, and authorship have the strongest relationship to post-release defects. The application of BN models helped explain the observed instability by demonstrating that the review-related predictors do *not* affect post-release defects directly and showed indirect effects. For example, changes that have *no review discussion* tend to be associated with files that have had many *prior defects* which in turn increase the number of post-release defects. We hope that similar analyses of other software engineering techniques may also yield a more nuanced view of their impact. Our replication package including our data and scripts is publicly available (Replication package 2018).

Communicated by: Tim Menzies

✉ Peter C. Rigby
peter.rigby@concordia.ca

Extended author information available on the last page of the article.

Keywords Code review measures · Statistical models · Bayesian networks

1 Introduction

For decades code review has been seen as a cornerstone of quality assurance for software projects. The process evolved from a formal process with checklists and face to face meetings (Fagan 2002) to a lightweight and semi-formal review done via e-mails or specially designed collaboration tools (Rigby and Storey 2011). The lightweight code review approach was originally used in open source software projects (OSS), because of their highly distributed nature (Mockus et al. 2000; Rigby et al. 2008) and has also become a common practice among commercial projects as well (Rigby and Bird 2013; Bacchelli and Expectations 2013). Recent studies suggest that the focus of review has shifted from early defect discovery to problem discussion and knowledge sharing (Bacchelli and Expectations 2013; Rigby and Bird 2013; Bosu et al. 2015; Rahman et al. 2017; Kononenko et al. 2016). It is perceived as a major quality control mechanism to prevent defects in production code (Bacchelli and Expectations 2013; Camilo et al. 2015; Munaiah et al. 2017).

An obvious and important scientific question is whether or not code reviews actually improve software quality, and whether our measurements of code review have explanatory power. To clarify such theoretical question, science resorts to replication to make it self-correcting system (Shull et al. 2008; Carver 1978). Replication helps establish if the phenomenon is dependable or idiosyncratic (Runeson and Höst 2009a; Shull et al. 2002; Axelrod 1997; ZCarver 2010). Our first aim is, therefore to conduct a similar-internal replication (a replication where only the experimenters varied (Gómez et al. 2014; Almqvist 2006)) of a highly reputable recent result investigating the effects code reviews have on software quality. We chose a commonly used quality measure: post-release defects. Such defects affect end-users (and vendor reputation) and are very costly to repair (Report 2002), and, therefore are a primary concern to software industry (Huang and Boehm 2006).

RQ 1. Replication: do previously reported associations between code review measures and post-release defects hold in a similar-internal replication study?

To investigate a hypothesis-driven scientific question researchers often use linear regression models¹ to examine the relation between code review (and other metrics) and software quality (Porter et al. 1998b; Kononenko et al. 2015; Morales et al. 2015). A recent award-winning work by McIntosh *et al.* (McIntosh et al. 2014; McIntosh et al. 2016) employed additive models to fit non-linear curves that are more suited for non-monotone or non-linear relationships than linear regression. We perform an exact replication of that experiment to determine if we can obtain the same conclusions using the same methods and data. Specifically, we construct OLS models with restricted cubical splines to model these relationships and discuss our findings with the first author, McIntosh, of this study to ensure that he agreed with our conclusions.

RQ 2. Differentiated-external replication: do previously reported associations between code review measures and post-release defects hold for another large software project?

Our second goal is to increase external validity (Gómez et al. 2014) of the results to avoid conclusions that are unique to the specific dataset reported in the paper. To accomplish this,

¹Machine learning methods focused on maximising prediction performance are widely used for defect prediction, but such methods are typically not transparent enough to test scientific hypothesis (Lin et al. 2013)

we apply exactly the same set of methods on a different software project: Chrome. This is sometimes referred as differentiated-external replication (Almqvist 2006). We chose the project due to its size and richness and quality of the associated data that allowed us to obtain measures highly similar to ones obtained in the Qt project of the replicated study. More specifically, we model software defects that are reported in a bug tracker. As control variables, we use many of the previously studied measures that have been shown to impact defects, including size, complexity, churn, authors, and file ownership (Bird et al. 2011; Hassan 2009). The focus of this study is on investigating code review measures many of which have been examined in past studies, including the number of reviewers, discussion length, and rushed review in a different setting (McIntosh et al. 2016; McIntosh et al. 2014; Kononenko et al. 2015; Rigby et al. 2008; Rigby et al. 2014).

RQ 3. Structure of the relationships: Are code review measures directly associated with post-release defects or are they affected by other measures of the development process that are, in turn, directly associated with post-release defects?

The findings from RQ1 and RQ2 point to the methodological limitations of linear regression and additive models when applied to datasets that have high correlations among the predictors as is typical software engineering data in general and in code review data in particular (Rigby et al. 2014; McIntosh et al. 2016). The linear (or additive) models can not reliably determine which of the highly correlated predictors are affecting the response. Principal Component Analysis (PCA) is typically applied in such cases but the results are hard to interpret because a linear combination of unrelated measures, e.g., combining lines of code, number of reviewers, and other unrelated concepts into a single predictor. This defeats the original purpose of testing the scientific hypothesis as discussed in Chapters 6.3, 6.7, and 10.2 of (James et al. 2013). Since automatic variable selection techniques are highly unstable (see, e.g. (Austin and Tu 2004)), best practices in empirical studies that employ regression models, recommend the manual removal of highly correlated variables, or variables that do not contribute to the explanatory power of the model. Such selection of variables relies on a subjective judgement of the researcher. Another shortcoming of such models is their inability to model the relations among predictors, which may reveal salient aspects of the development process by providing a rich picture of how the predictors may influence each other and the response.

A Bayesian Network (BN) is a Probabilistic Graphical Model (PGM). PGM describes probabilistic relationships among variables that describe a problem domain (Heckerman 1998). This model has several advantages over linear or additive regression models. In particular, it allows for a natural representation of conditional dependence and independence using graph notation where variables are nodes and dependencies are edges. The removal of the notion of predictor and response variables disposes of the oversimplifying assumption that a single response variable is explained by a long list of predictors. Instead the edges in the Bayesian Network provide a meaningful structure based on collected data. Each variable in a graph can be interpreted as a predictor or a response variable based on the topology of the graph. The researcher can then inject information to understand the impact of an edge of interest (Friedman et al. 1999).

Our main findings from RQ 1, the reproduction of the study by McIntosh *et al.* (McIntosh et al. 2016), have demonstrated high sensitivity of the regression modeling results to the subjective steps in the analyses when data contains highly correlated predictors. In particular, we found that even in exact reproduction we were unable to confirm the predictive power of code review measures on post-release defects. We discussed the finding with McIntosh and, according to his opinion, code review measures are not likely to explain more of the variance than traditional measures. Moreover, the results are inconsistent across software

releases and heavily depend on the variables selected. We did, however, find several metrics not related to code reviews, such as churn or prior defects, that were reproduced reliably despite the subjectivity of the variable selection process.

The investigation in RQ 2, increased the external validity of the findings by confirming that a relatively small set of measures, churn and prior defects, are related to post-release defects on a large and unrelated software project. As on the Qt dataset, the impact of review measures was inconsistent across software releases and heavily depend on the variables selected.

In RQ 3, to reduce the subjectivity of variable selection process and to untangle the complex web of dependencies among the predictors, we applied Bayesian Networks (BN) on both datasets. The approach revealed that there is no direct relation between review measures and defects. The graph shows, for example, that modules with more self-approved changes also have more changes with no discussion, more reviewers, and also more prior defects. An increase in review issues increases the share of the work done by the minor authors, which, in turn, is associated with increased number of defects.

This paper is organized as follows. In Section 2, we discuss the case study design, the systems under study, and the data extraction process. We also give a brief overview of the Chrome code review process. In Section 3, we replicate and reproduce McIntosh *et al.*'s (McIntosh et al. 2016) study, describe the model construction, results, and discussion. In Section 4, we describe BNs and discuss the findings from these models. Threats to validity are discussed in Section 5. The final section concludes the paper and suggests future work. Our replication package including our data and scripts is publicly available (Replication package 2018).

2 Case Study Design and Data

In this section we discuss the case study design including the projects under study and reasons for their selection. We describe the data sources, steps in the data extraction, and analysis approach. We discuss the Bayesian Network modeling methodology in Section 4.

2.1 Systems Under Study

McIntosh et al. (McIntosh et al. 2016) mined code review data from Android, LibreOffice, QT, ITK, and VTK. They did not conduct an analysis on Android and LibreOffice because they found that many of the reviews were not linked to bug reports which did not allow them to study the impact of review on bugs. In total, they studied two QT releases and one release for VTK and ITK. For the reproduction, McIntosh provided the Qt and ITK data that was used in their work (McIntosh et al. 2016). The ITK data had only 24 defective components and 344 commits with reviews. We feel that this dataset is too small to produce meaningful statistical models. Although we include the ITK results in our replication package (Replication package 2018), we only present the Qt results in this work. To improve external validity, we replicate the study on the Google Chrome project, because like QT, it is large and primarily written in C++. A further reason for studying Chrome is that it is an open source web browser, that is mostly developed by paid Google developers and its development practices mirror those used internally at Google. Chrome developers are required to perform code review on each change and use Reivfield, the precursor to Gerrit, to improve traceability of bugs, changes, and reviews.

For completeness, we briefly describe Chrome’s code review process which resembles other modern review practices (Rigby et al. 2012). A review begins when the change author submits a patch and invites reviewers. A reviewer examines a change and either approves it by replying with special keyword *lgtm* (looks good to me) or proposes improvements. The author addresses comments either by fixing issues in code or by replying to the reviewer comments. Subsequent modifications to the original patch appear in the same review and are called *patchsets*. The new patchset triggers a new cycle of review and revision. The process continues until all issues are fixed and the reviewers are satisfied with the patch. The code can then be merged to the trunk.

2.2 Chrome Data Extraction

To understand the influence of code review measures on post release defects we need to create a link between the code review, the source files, and reported bugs. We collect data from three data sources: Reitviold, Git, and the Chrome bug tracker (Fig. 1). The data extraction is divided into the three steps described below.

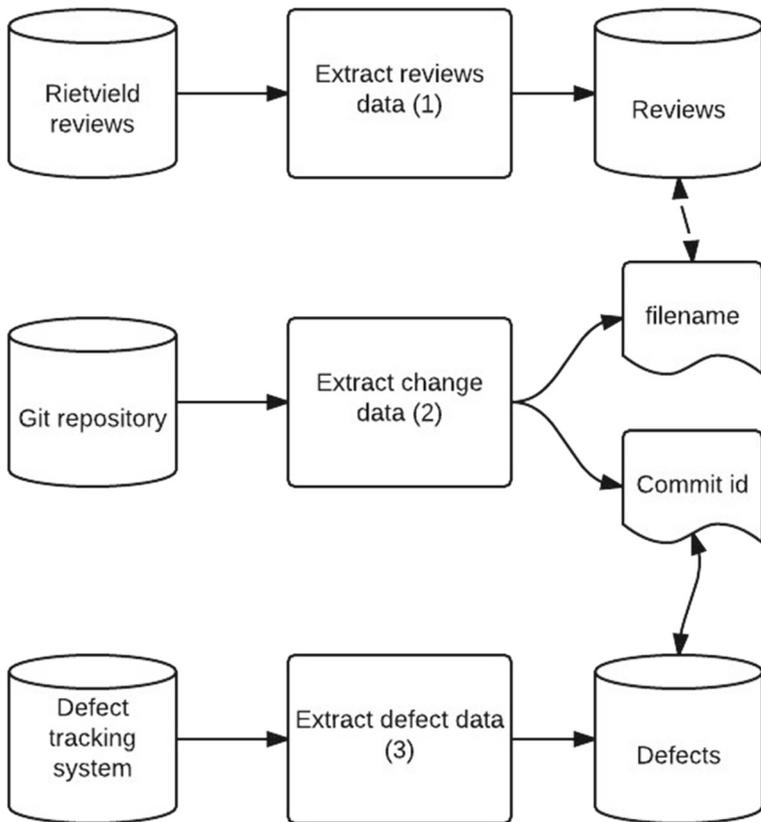


Fig. 1 Chrome data sources and extraction methodology. Reviews are extracted from the Rietviold review system. Bugs are extracted from the defect tracking system. The bug and review ids are contained in the Git commit and linked to the modified files and directories

Extracting review data: We use the Reitvield API to download code reviews in JSON format and extract the data into a database. For each code review patch revision we extract the unique identifier and the set of files modified by this revision. For every file and revision we also capture the number of added and removed lines to calculate the size of a change. We process the reviewers comments. We ignore comments that were added automatically by a bot or by the patch author.

Extracting Git repository information: We extract commit information *i.e.* the commit hash and list of files related to the change from the Git repository. We use the Understand static analysis toolkit² to extract source code measures from the files.

Extracting defect data: We mine the defects from the Chrome issue tracker by scraping the pages. We extract the submission date, type of the issue, review ID, and commit ID for the fix.

Post-release defects: We consider a defect to be the post-release defect of the current release if it was submitted during the time period between the release dates of the current and the following releases. We use Chrome release calendar website for release dates information.³ Following McIntosh (Mcintosh et al. 2016), we associate the post-release defects with the pre-release reviews and other source measures using first the file level and then sum the measures to the component, *i.e.* directory level. The directory was chosen as the unit of analysis to reduce the fraction of zero observations because the majority of the files in the system do not have any defects.

2.3 Collected Measures

The measures we collect to evaluate the impact of code review on post-release defects are well known and have been used in multiple past studies (Mcintosh et al. 2016; McIntosh et al. 2014; Rigby and Bird 2013; Kononenko et al. 2015) and are described in Table 1. We divide them into four categories: product, process, human factors, review participation, and reviewer expertise. The code review measures are the number of reviewers, discussion length, rushed reviews, typical reviewer expertise, etc. The control variables in our model are also well known and widely used with defect prediction models (Mockus et al. 2000; Hassan 2009; Bird et al. 2011; Graves et al. 2000). The control variables include the size of the file, the number of prior defects, and the churn.

3 Code Review Replication and Reproduction Study

We replicate the study published by McIntosh et al. (Mcintosh et al. 2016). We strictly follow the steps of the model construction described in the original paper (Mcintosh et al. 2016). We fit an Ordinary Least Squares (OLS) regression model. Since the dependent variable is the number of post-release defects and it is highly skewed, we log transform it. We also create regression models that take non-linear effects into account. We then compare the goodness of fit among models and discuss the contribution of each independent variable.

Correlated variables can distort the contribution of a variable to a model and must be removed. We use the hierarchical clustering analysis (Figs. 2 and 3) with the threshold of $|\rho| \geq 0.7$ suggested by McIntosh (Mcintosh et al. 2016) to identify the highly correlated

²<https://scitools.com/>

³<https://www.chromium.org/developers/calendar>

Table 1 Description of Measures: product, process, human factors, review participation, and reviewer expertise

	Measure	Description
Product	Size	Number of lines of executable code in component
	Complexity	The McCabe cyclomatic complexity.
Process	Prior defects	Number of defects fixed in component prior to the considered release period
	Effective tests (Chrome only)	Total number of times a test found an issue during the review process
	Churn	Sum of added or removed lines of code per component during considered period of time
	Change entropy	Distribution of changes among files within a component
Human Factors	Minor authors	Number of unique contributors that contribute less than 5% of code changes to a component
	Major authors	Number of unique contributors that contribute at least 5% of code changes to component
	All authors	Number of unique contributors to component
	Author ownership	Proportion of changes to component done by major authors
Review Participation	Rushed reviews	Number of reviews that were concluded faster than acceptable review rate (200 loc per hour)
	Changes without discussion	Changes that were integrated without discussion comments
	Self approved changes	Changes that were approved for integration only by the author
	Typical discussion length	Discussion length typical for that specific component measured in number of discussion comments. Normalized by size of change (churn)
	Typical review window	The amount of time between the patch upload and its approval for integration. Normalized by size of change (churn)
	All reviews	Total number of times the component was reviewed
Review Expertise	All reviewers	Total number of reviewers that reviewed a component
	Review issues	Total number of patch revisions created during review process
	Effective reviews (Chrome only)	Number of revisions that led to a code change during a single review per component
	Lacking subject matter expertise	Number of changes that were not authored or approved by major author
	Typical reviewer expertise	Total number of changes to the component authored or reviewed by this reviewer prior to this change

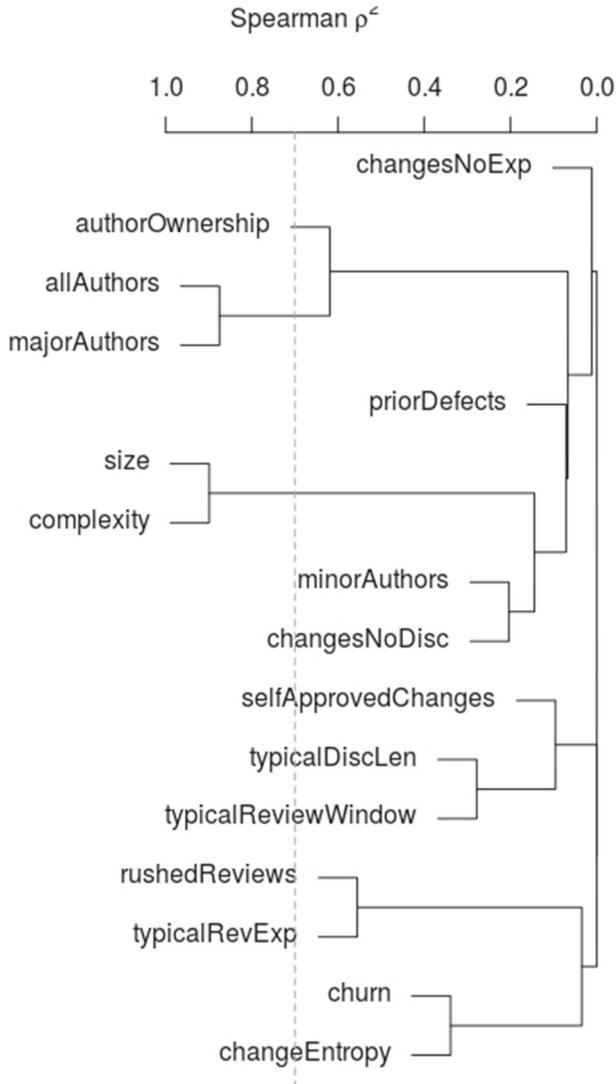


Fig. 2 Hierarchical Correlation Analysis for Qt 5.0. For variables correlated at $|\rho| \geq 0.7$ the simpler measure is kept. We also conduct a redundancy analysis

variables. Then we select which variables will be discarded using drop-one analysis and the principle of parsimony. The results of this step are summarized in the Tables 2 and 3.

Redundant variables can be explained by the remaining variables, that is they do not contribute to the explanatory power of the model and should be removed. Such variables may be overlooked by the pairwise correlation analysis, therefore we use *redun* function from *Hmisc* R package.⁴ For each independent variable a regression model is fitted using as predictors the remaining variables. If the model has a R^2 greater than 0.9, then the current

⁴<https://cran.r-project.org/web/packages/Hmisc/Hmisc.pdf>

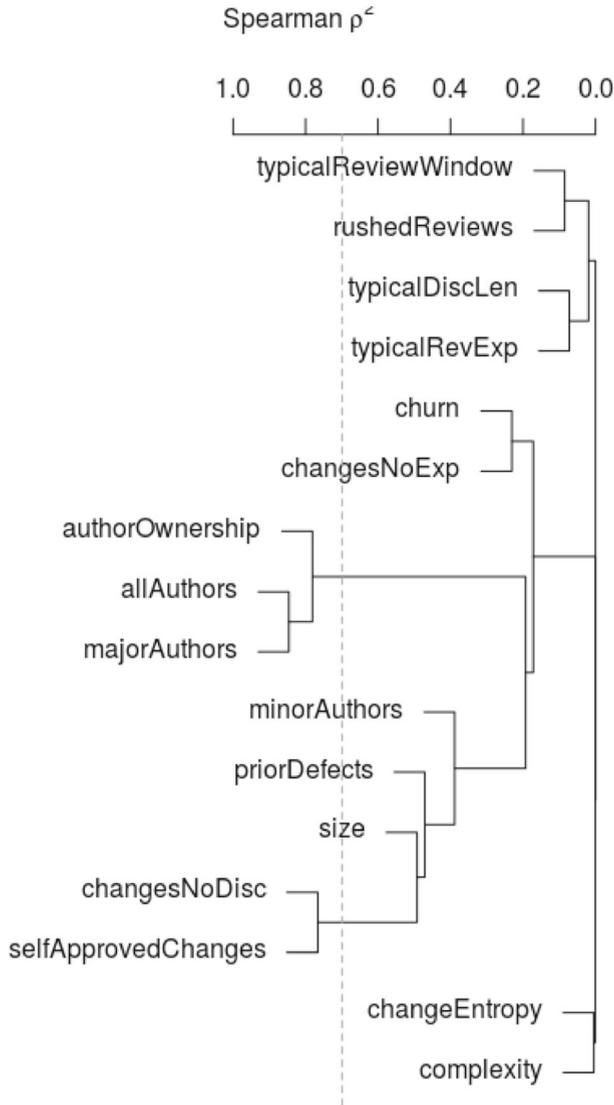


Fig. 3 Hierarchical Correlation Analysis for Chrome 40. For variables correlated at $|\rho| \geq 0.7$ the simpler measure is kept. We also conduct a redundancy analysis

variable is considered to be redundant because the linear combination of the remaining variables can closely approximate this variable.

Non-linear effects and degrees of freedom. Traditional defect prediction models assume linear dependencies between the dependant and independent variables. McIntosh et al. (McIntosh et al. 2016) showed that for some code review measures this relation has a non-linear shape. To identify variables that may be nonlinear, we calculate the Spearman multiple ρ^2 scores for each independent variable. Variables with higher scores are more likely to be non-linear. To fit a non-linear curve we use restricted cubic splines in the *rms* R package (Harrell Jr. 2013). Using this approach we assign knots, which are points where

Table 2 Post-release defects prediction model for Qt. Original and replication study results. Non-linear models do not outperform linear models. While statistically significant, review measures are unstable and have little predictive power compared to traditional measures

Release	5.0 McIntosh	5.0	5.1 McIntosh	5.1
Nonlinear model adjusted R^2	0.69	0.62	0.46	0.66
Linear model adjusted R^2		0.61		0.63
Nonlinear model w/o codereview variables adjusted R^2		0.61		0.64
Overall D.F.	80	81	78	81
Allocated D.F.	22	22	24	22
	Overall	Nonlinear	Overall	Nonlinear
Size	4	3	2	1
D.F.	110***	76***	10**	5*
χ^2	1	na	1	na
Complexity	1°	na	< 1°	na
D.F.	‡	‡	2	1
χ^2	‡	‡	48***	45***
Prior defects	1	na	1	na
D.F.	1°	na	< 1°	na
χ^2	2	1	2	1
Churn	8*	7**	6*	6*
D.F.	‡	‡	3	2
χ^2	‡	‡	274***	63***
Change entropy	‡	‡	1	na
D.F.	‡	‡	2°	na
χ^2	‡	‡	193***	13***
All authors	‡	‡	1	na
D.F.	‡	‡	2°	na
χ^2	‡	‡	‡	‡
Minor authors	‡	‡	‡	‡
D.F.	‡	‡	‡	‡
χ^2	‡	‡	‡	‡

Table 2 (continued)

	5.0	5.0 McIntosh	5.0	5.1 McIntosh	5.1
Release					
Nonlinear model adjusted R^2	0.62	0.69	0.62	0.46	0.66
Linear model adjusted R^2	0.61		0.61		0.63
Nonlinear model w/o codereview variables adjusted R^2	0.61		0.61		0.64
Overall D.F.	81	80	81	78	81
Allocated D.F.	22	22	22	24	22
Major authors	†	†	†	†	†
χ^2	†	†	†	†	†
D.F.	†	†	†	†	†
Author ownership	†	†	†	†	†
χ^2	†	†	†	†	†
D.F.	1	1	na	1	1
Self-approved	7*	22***	1°	< 1°	< 1°
χ^2	2	†	†	na	na
D.F.	2	†	†	1	2
Rushed reviews	4°	†	†	48***	3°
χ^2	2	2	2	23***	3°
D.F.	1	1	1	1	2
Changes w/o disc.	18**	6°	4*	3°	36***
χ^2	1	†	†	†	1
D.F.	< 1°	†	†	†	< 1°
Typical review window	1	†	†	†	na
χ^2	1	†	†	†	na
D.F.	1	3	3	2	1
Typical disc. length	< 1°	26***	24***	32***	3°
χ^2	na	†	†	21**	na

Table 2 (continued)

		5.0 McIntosh	5.0	5.1 McIntosh	5.1
Release					
Nonlinear model adjusted R^2		0.69	0.62	0.46	0.66
Linear model adjusted R^2			0.61		0.63
Nonlinear model w/o codereview variables adjusted R^2			0.61		0.64
Overall D.F.		80	81	78	81
Allocated D.F.		22	22	24	22
Lacking subject matter expertise	D.F.	2	2	1	1
	χ^2	80***	70***	34***	22**
Typical reviewer expertise	D.F.	4	3	2	1
	χ^2	26***	24***	32***	21**
Changes w/o disc.	D.F.	2	2	2	2
	χ^2	6°	18**	3°	36***
Typical review window	D.F.	†	†	†	†
	χ^2	†	< 1°	†	< 1°
Typical disc. length	D.F.	4	3	2	1
	χ^2	26***	24***	32***	21**
Lacking subject matter expertise	D.F.	2	2	4	3
	χ^2	80***	70***	34***	22**
Typical reviewer expertise	D.F.	4	3	2	1
	χ^2	26***	24***	32***	21**

Discarded during: † - Removed during correlation analysis; ‡ - Removed during redundancy analysis
 Statistical significance: ** * $\rho < 0.001$; * * $\rho < 0.01$; * $\rho < 0.05$; ° $\rho > 0.05$
 Other: na - not used

Table 3 Post-release defects prediction model for Chrome. Non-linear models do not outperform linear models. While statistically significant, review measures are unstable and have little predictive power compared to traditional measures

Release	39	40	41	42	43	44
Nonlinear model adjusted R^2	0.61	0.58	0.59	0.59	0.51	0.53
Linear model adjusted R^2	0.59	0.54	0.56	0.57	0.50	0.49
w/o coder review variables adjusted R^2	0.60	0.58	0.59	0.59	0.49	0.53
Overall D.F.	62	87	90	84	83	80
Allocated D.F.	26	21	23	20	19	20

	Overall	Nonlinear	Overall	Nonlinear	Overall	Nonlinear	Overall	Nonlinear	Overall	Nonlinear
Size	D.F.	4	3	2	1	4	3	2	1	2
	χ^2	28***	3°	1°	1°	10*	< 1°	1°	< 1°	3°
Complexity	D.F.	1	na	1	na	1	na	1	na	1
	χ^2	< 1°	na	3°	na	< 1°	na	< 1°	na	< 1°
Prior defects	D.F.	4	3	2	1	4	3	4	3	2
	χ^2	43***	42***	37***	34***	74***	71***	61***	55***	21***
Churn	D.F.	2	1	2	1	2	1	2	1	2
	χ^2	14**	8*	23***	22***	31***	30***	26***	24***	< 1°
Change entropy	D.F.	2	1	1	na	1	na	1	na	1
	χ^2	< 1°	< 1°	< 1°	na	< 1°	na	< 1°	na	< 1°
All authors	D.F.	3	2	4	3	4	3	3	2	3
	χ^2	49***	2°	43***	8*	42***	2°	33***	< 1°	51***
Minor authors	D.F.	†	†	†	†	†	†	†	†	†
	χ^2	†	†	†	†	†	†	†	†	†
Major authors	D.F.	†	†	†	†	†	†	†	†	†
	χ^2	†	†	†	†	†	†	†	†	†
Author ownership	D.F.	†	†	†	†	†	†	†	†	†
	χ^2	†	†	†	†	†	†	†	†	†

Table 3 (continued)

Release	39	40	41	42	43	44
Nonlinear model adjusted R^2	0.61	0.58	0.59	0.59	0.51	0.53
Linear model adjusted R^2	0.59	0.54	0.56	0.57	0.50	0.49
w/o coder review variables adjusted R^2	0.60	0.58	0.59	0.59	0.49	0.53
Overall D.F.	62	87	90	84	83	80
Allocated D.F.	26	21	23	20	19	20
	Overall	Nonlinear	Overall	Nonlinear	Overall	Nonlinear
Self-approved	D.F. 4	3	4	3	3	3
	χ^2 8*	7°	12*	5°	2°	3°
Rushed reviews	D.F. 1	na	1	na	1	na
	χ^2 1°	na	19***	na	na	na
Changes w/o disc.	D.F. †	†	†	†	†	†
	χ^2 †	†	†	†	†	†
Typical review window	D.F. 1	na	1	na	1	na
	χ^2 1°	na	1°	na	na	na
Typical disc. length	D.F. 1	na	1	na	1	na
	χ^2 < 1°	na	< 1°	na	< 1°	na
Lacking subject matter expertise	D.F. 2	1	2	1	1	1
	χ^2 7*	3°	32***	7***	5*	10***
Typical reviewer expertise	D.F. 1	na	1	na	1	na
	χ^2 < 1°	na	< 1°	na	na	na

Discarded during: † - Removed during correlation analysis; ‡ - Removed during redundancy analysis

Statistical significance: *** ρ < 0.001; ** ρ < 0.01; * ρ < 0.05; $\circ\rho$ \geq 0.05

Other: na - not used

the slope changes, to potentially non-linear variables. The more knots that are added the greater the curve complexity. Every additional knot requires a degree of freedom. If we use all of the degrees of freedom, then there will be a knot for each data point and the fit will be perfect, but the model will be overfitted to the data. As a result, the degrees of freedom are budgeted to avoid over fitting while still allowing variables that have a high ρ^2 score to be modelled non-linearly.

To *assess model fitness*, we report the adjusted R^2 to compensate for the large number of variables (McIntosh et al. 2016). To assess the individual contributions of each variable, we report its statistical significance and the Wald χ^2 maximum likelihood test value. The larger the value the greater the impact the variable has on the model. The results are summarized in Tables 2 and 3. The tables also contain the results from McIntosh et al. (McIntosh et al. 2016) and conform to the same structure.

3.1 Variable Selection and Model Construction

In our summary table we have approximately 1.3k reviews for Qt 5.0 and 1.4K for Chrome 40. We start with 16 measures. This gives us 81 and 87 degrees of freedom for Qt and Chrome respectively. Following previous works, we discard measures with correlation at or above 0.7. We use clustering analysis to identify these measures, see Figures 2 and 3. We also perform *drop one* analysis to determine which measures should be discarded from each cluster. *Major authors*, *author ownership* and *complexity* were removed from Qt dataset. *Major authors*, *author ownership* and *changes without discussion* were removed from Chrome. Using a redundancy test *minor authors* was removed from both datasets. We perform a non-linearity analysis. Variables that exhibit a higher degree of non-linearity require additional degrees of freedom to model their curved line. The results of the variable selection process and the number of allocated degrees for each variable can be found in in Tables 2 and 3.

To represent our models, we use the R language notation (Pinheiro et al. 2011). For example, the formula $y \sim a + b$ means that the response y is modelled by explanatory variables a and b . McIntosh (McIntosh et al. 2016) used the following model for Qt:

$$\begin{aligned} \log(\text{defects} + 1) &\sim \text{rcs}(\text{size}, 5) + \text{rcs}(\text{all authors}, 5) \\ &+ \text{complexity} + \text{churn} + \text{rcs}(\text{change entropy}, 3) \\ &+ \text{rcs}(\text{changes w/o discussion}, 3) \\ &+ \text{rcs}(\text{self} - \text{approved changes}, 5) \\ &+ \text{rcs}(\text{typcal discussion length}, 5) \\ &+ \text{rcs}(\text{typical reviewer expertise}, 5) + \text{rcs}(\text{lacking subject matter expertise}, 5) \end{aligned}$$

Our Qt model is the following:

$$\begin{aligned} \log(\text{defects} + 1) &\sim \text{rcs}(\text{size}, 5) + \text{rcs}(\text{prior defects}, 5) \\ &+ \text{rcs}(\text{churn}, 3) + \text{rcs}(\text{changeentropy}, 3) \\ &+ \text{rcs}(\text{all authors}, 5) + \text{rcs}(\text{changes w/o discussion}, 5) \\ &+ \text{self} - \text{approved changes} + \text{typcal discussion length} \\ &+ \text{typcal review window} + \text{rcs}(\text{rushed reviews}, 3) \\ &+ \text{rcs}(\text{lacking subject matter expertise}, 3) + \text{typical reviewer expertise} \end{aligned}$$

Our Chrome model is the following:

$$\begin{aligned} \log(\text{defects} + 1) \sim & \text{rcs}(\text{size}, 5) + \text{rcs}(\text{prior defects}, 5) \\ & + \text{complexity} + \text{rcs}(\text{churn}, 3) + \text{rcs}(\text{change entropy}, 3) \\ & + \text{rcs}(\text{all authors}, 5) + \text{rcs}(\text{self - approved changes}, 5) \\ & + \text{typical discussion length} + \text{rushed reviews} \\ & + \text{typical review window} + \text{rcs}(\text{lacking subject matter expertise}, 3) \\ & + \text{typical reviewer expertise} \end{aligned}$$

Restricted Cubic Splines are represented by the *rcs* function in the formula where the first argument is the predictor and the second argument is the number of knots (the amount of nonlinearity) permitted. We fit OLS model using formulas from above and calculate adjusted R^2 to assess goodness of fit. Tables 2 and 3 summarize the results and contain the original results from McIntosh et al. (McIntosh et al. 2016). In summary, the response variable is modeled via linear and non-linear dependencies approximated via cubic splines.

3.2 Model Results and Model Comparisons

In this section we compare our replication results with those from McIntosh et al. (McIntosh et al. 2016) original study and new results from Chrome. We highlight differences and discuss their possible causes.

3.3 Comparing Linear and Non-Linear Models

To illustrate the nonlinear effect we select an independent variable with the highest potential of nonlinearity from the model and calculate predicted number of post-release defects as the function of this variable, using *Predict* function from R rms package. The rest of the variables are fixed at their median values. As an illustration, we choose *prior defects* for Qt because it had the highest Spearman squared value among independent variables participating in the model. We then plot the results in Figure 4. Although the shape of the plot may suggest some nonlinearity, the grey funnel, which is the error margin, is too wide to claim with confidence that these variables have a nonlinear relation with the response variable. The goodness of fit R^2 also shows that nonlinear models do not yield better results than regular linear models. Discussion with McIntosh revealed that the text of the original paper was ambiguous and they only log transformed the dependent variable. We find that a log transformation of skewed independent variables provides a reasonable model without adding the complexity of a non-linear model.

3.4 Models with and without Review Measures

The results show that although many of the code review measures are statistically significant, they usually tend to have lower values of Wald χ^2 test than other measures, suggesting their lower contribution to explanatory power of the model. Even the most prominent measures, like *typical discussion length* and *rushed reviews* are repeatedly outperformed by measures like *size*, *prior defects*, and *all authors*. As a further investigation, we fit a model *without* review measures and record the values of adjusted R^2 (shown in bold in the section of R^2 values in the Tables 2, 3). The decrease in the values of adjusted R^2 in both datasets is minimal, meaning that overall contribution of the review measures to the explanatory power of the model is limited. In addition to low contribution to the model the performance

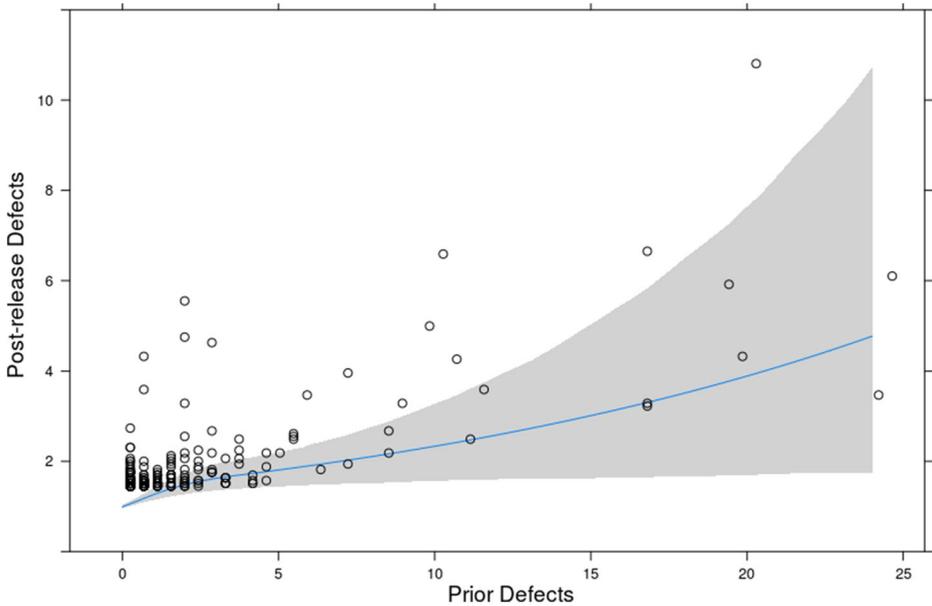


Fig. 4 A wide margin of error for nonlinear predictions, for example, prior defects in Qt 5.0. Nonlinear models are unnecessary, see Tables 2 and 3

of review variables is inconsistent between datasets and releases. For instance, in McIntosh et al. *rushed reviews* was discarded from the model in Qt 5.0 during the correlation analysis, however, in Qt 5.1 this measure is one of the strongest predictors of post-release defects. The *typical discussion length* is one of the most influential variables for both Qt releases in both studies, but in Chrome dataset the contribution of this variable is insignificant. Discussing our results with McIntosh, he stated that he did not believe that review measures could dominate traditional measures. Our answer to RQ1, Replication, is summarized below.

Conclusion 1: *The review measures contributed little to the performance of the model, with the R^2 remaining almost unchanged from the model that included only the traditional predictors, such as the number of prior-defects, size, and authors.*

3.5 Impact of Individual Variables

Size of component is a well-known predictor in empirical software studies. McIntosh et al. show that in the Qt project *size* provides significant contribution to the explanatory power of the model. Our result is similar for Qt dataset. However, in Chrome dataset the contribution of the *size* measure is quite small (Table 3).

Prior defects and all authors have been shown to be good predictors of future defects (Graves et al. 2000; Bird et al. 2011). McIntosh et al. discard *prior defects* in Qt 5.0 due to redundancy. In our study, the redundancy analysis on the Qt 5.0 dataset does not indicate that *prior defects* are redundant and, on the contrary, is statistically significant. For the Qt 5.1 release, both McIntosh et al. and our model keep *prior defects* but find it to be a poor predictor. The *all authors* measure is redundant in Qt 5.0 release in our study contrary

to the McIntosh et al.. For Qt 5.1 the *all authors* is the most influential predictor in the model. This result is replicated in both studies. In Chrome dataset these two variables are repeatedly found to be the most influential variables of the model. A possible explanation for this inconsistency could be that these two variables share a common cause. Defects are not always fixed by the owner of the module, especially in big teams. That means that more developers are touching the file, and the more developers modifying a file the higher the risk of the future defects. Intuitively, the growth in these two measures should be related, but our correlation and redundancy analysis fails to find this. These inconsistent results suggest that traditional variable selection techniques are not capable of coping with the high correlations in our datasets and indicate the need for a different approach that can deal with complex interactions between variables.

Review measures . The important measures in the Qt dataset are similar to what McIntosh et al. found. The *self-approved changes* has low impact on post-release defects. The *rushed reviews* variable was discarded in 5.0 release, but in 5.1 it appears as one of the most influential variables. The *typical discussion length* variable has moderate to strong influence in both releases. For the Chrome dataset the review measures are not statistically significant in most cases. When they are, such as in *lacking subject matter expertise* the result is inconsistent across releases. The overall performance of review variables is inconsistent in both studies suggesting the following conclusion to RQ 2, Differentiated-external replication.

Conclusion 2: *The inconsistent performance across projects and releases of strong predictors, like all authors, prior defects and others, suggests a possible issue with the traditional variable selection approach and indicates the need for an approach that is capable of dealing with a more complex inter-variable relations.*

4 Bayesian Networks Models

To address the concern of the unexpected absence of the relationship between code review measures and post-release defects demonstrated in the previous models, and the lack of reproducibility due to the subjectivity in variable selection approaches that are necessary in a traditional model, such as the one used in Section 3, we use Bayesian Networks (BN) as an alternative modeling approach. Our goal in using the BN model is not to create the best predictive model for post-release defects. Instead, we focus on understanding the complex interaction between the variables described by the data, and determining which variables directly impact the number of post-release defects in such a generative model.⁵

4.1 Background: Bayesian Network

Bayesian Network models have several advantages over regression models. To be precise, regression analysis is a very simple BN where there is one directed link from each independent variable to the dependent variable. BNs, therefore, can help with multicollinearity

⁵A generative model specifies a joint probability distribution over all observed variables, whereas a discriminative model provides a model only for the target variable(s) conditional on the predictor variables. Thus, while a discriminative model allows only sampling of the target variables conditional on the predictors, a generative model can be used, for example, to simulate (*i.e.* generate) values of any variable in the model, and consequently, to gain an understanding of the underlying mechanics of a system, generative models are essential.

by establishing the relationships among independent variables. In the process of BN construction we can control the number of edges (relations) by specifying a connection strength threshold. Once the Bayesian Network is constructed we can use the graphical representation to learn about less obvious interactions among variables and infer how the injection of specific facts affects variables of interest. We use BN to investigate the lack of consistency in the replication in the previous sections.

One important concept related to the BNs is the concept of *Markov Blanket* (Pearl 2014). The Markov Blanket for a node in a Bayesian Network is the set of nodes composed of its parents, its children, and its children's other parents (co-parents). The Markov blanket of a node contains all the variables that shield the node from the rest of the network *i.e.* for a node A , its Markov Blanket MB_A , and a node $B : B \neq A, B \notin MB_A$, we have the property that:

$$Pr(A|MB_A, B) = Pr(A|MB_A)$$

This means that the Markov blanket of a node is the only knowledge needed to predict the behavior of that node.

We will construct a BN model without assuming any domain expertise, edges, or assumptions about prior data distributions. Instead we will use minimal a-priori model that focuses on the search for the best Bayesian graphical representation for the dataset (*i.e.* structure learning using hill climbing).

Despite the promises of BNs, they tend to be quite sensitive to data, and operational data is often problematic (Mockus 2014; Zheng et al. 2015). Careful preprocessing is needed to ensure a reliable and reproducible result. We next discuss discretization of variables and structure learning with hill climbing approaches used to address these concerns.

4.1.1 Discretization

For our regression model, we found that all our variables have a long-tailed distribution that could not be corrected even by a log-transformation. Since BN structure learning methods for continuous data require a normal distribution, we discretize the data, as is often done with prediction that involves classifiers (Scutari 2013). Discretizing variables while preserving relationships among them is an NP-hard problem (Chlebus and Nguyen 1998), but several heuristics exist. The commonly used supervised methods optimize discretization to improve explanatory power for a single response variable, such as, Chi-square, or MDLP. However, these are not suitable for a BN structure search, because we do not know *a-priori* which variables will be responses (have arrows pointing to them) and which will be independent (have no incoming arrows). While some research on multidimensional discretization methods exist (Perez et al. 2006), we are not aware of any such method that has a robust implementation in a statistical package. We, therefore, use unsupervised discretization methods. The added benefit is that the discretization was totally response-variable agnostic unlike the commonly used supervised discretization methods, which prevents any bias towards specific fit that may accompany supervised methods.

Following the recommendations from Garcia *et al.*'s (Garcia et al. 2013) survey, we use the Equal Frequency discretization method and the implementation in the *arules* package. The *defects* node was discretized to a binary no-defect/defect variable, because around 73% of the directories have no defects, therefore it makes sense to just predict whether or not there will be a defect for our dataset. Two levels were also assigned to minor authors, rushed reviews, typical review window, and lacking subject matter expertise because more than 50% of entries were zero. Based on the data distributions for the remaining variables three

levels were deemed appropriate. We present the distribution of the variables in our replication package (Replication package 2018) as additional evidence for the choice of our discretization levels.

4.1.2 BN Structure Search: Hill Climbing

To learn the BN structure from our data, we chose a well-performing and widely used (Dey and Mockus 2020) Hill-Climbing (HC) algorithm from the *bnlearn* R package. This HC algorithm attempts to maximize the network score with several scoring functions available in the *bnlearn* package: *e.g.*, BIC, AIC, BDE. A detailed study examining how well different scores performed concluded that in general all scores perform similarly and for large data sets Bayesian scores are more suitable (Carvalho 2009). Since our dataset is not particularly large, at least for the individual releases, we decided not to use Bayesian scores *e.g.*, BDE, instead we chose to focus on the information theoretic scores *e.g.*, AIC, BIC. We finally selected the BIC score because it is more appropriate for constructing explanatory models, while AIC is better suited for building predictive models (Sober 2002; Shmueli 2010).

Hill-Climbing has the known limitation of finding a local maxima, and there are several enhanced versions of the algorithm that deal with this shortcoming. The R implementation provides parameters for the number of random restarts and perturbations as tuning parameters to deal with this problem. However, these parameters can make the results noisy, with different settings inducing slightly different networks. To mitigate this effect, we use the non-parametric bootstrap model averaging method, which provides confidence levels for both the existence of an edge and its direction (Friedman et al. 1999). This enables us to select a model based on a confidence threshold. Friedman et al. (Friedman et al. 1999) argued that the threshold is domain specific and needs to be determined for each domain. To identify a suitable threshold, we performed a simulation study, by generating a simulated dataset for the same number of nodes. The result of the simulation showed that a threshold of 0.65 was suitable to accurately recover the original structure. We also investigated alternative thresholds to assess the stability of the results as described in Section 5.

Finally, due to the HC algorithm not being a deterministic one, we repeated the process of generating a model 100 times, according to the recommendation by Arcuri and Briand (Arcuri and Briand 2011) and generated our final model based on the averaged result of these 100 runs.

4.1.3 Combining Data for Qt and Chrome Releases

We created new datasets by combining the data for all the Qt releases, the Chrome releases, and also combined the data for all releases for the two projects, which gave us three new datasets. Combining the datasets makes our final model more robust and the result is arguably more generalizable, because of having more training examples and not being prone to overfit to the specific characteristics of one particular release.

4.1.4 BN Model Performance Measures

Model performance can be evaluated using explanatory and predictive performance measures (Shmueli 2010). We first create an explanatory model that allows us to understand which software engineering measures have the greatest influence on post-release defects. We use the variance explained, which is the proportion of the log-likelihood score of the model relative to the baseline model which assumes that all the variables are independent.

We test the predictive power of our BN models for the purpose of comparison between similar models and to demonstrate the practical applicability of the models. Training and testing the models with Cross-fold validation is *not* appropriate because we have time ordered data. To ensure that we are not using future data to predict past observations, we trained our model on the earlier 70% of the data and test it on the subsequent 30% of the data. We use the *Accuracy* and Cohen's *Kappa* measures as the performance measures for our models.

4.1.5 CPT: the probability of having a Defect Given each Variable

To determine the individual impact that each variable has on post-release defects, we create Conditional Probability Tables (CPT) for each variable in the Markov Blanket that directly influences defects. The CPTs are trained with the *gRain* package in R, using Junction Tree belief propagation method (Lauritzen-Spiegelhalter algorithm (Højsgaard et al. 2012; Lauritzen and Spiegelhalter 1988; Koller and Friedman 2009)) from the BN models. See (Højsgaard et al. 2012) for details about how the CPT tables are calculated. The code for CPT construction is available in our replication package (Replication package 2018).

4.2 Results for BN Models

We created BN models for each release, each project, and for the entire combined dataset. The results are summarized in Table 4. Importantly, the post-release defects variable, *i.e.* the *defects* node in the Markov Blanket, was influenced by but did not influence any of the other variables. This is reassuring, because it is not possible for post-release defects to influence pre-release measures (going backward in time). Our models were able to properly account for the fact by correctly identifying the prior defects in the *priordefects* variable as having the influence on post-release defects.

Table 4 shows the model performance measures: variance explained, Accuracy, and Kappa. We observe that Kappa values varied between 0.41 and 0.51, which signifies a fair or moderate agreement according to both Landis and Koch (Landis and Koch 1977) and Fleiss (Fleiss et al. 1981). Similarly, the accuracy of the models was also observed to be

Table 4 Variables in Markov Blanket of *defects* for each BN model and the model performance as measured by Variance Explained, Accuracy, and Kappa

Release	Variables in Markov Blanket of <i>defects</i>	Variance Explained	Accuracy	Kappa
Qt_50	priordefects, changesnodisc	51.4%	0.85	0.43
Qt_51	allauthors, size	32.9%	0.86	0.44
Qt_combined	priordefects, changesnodisc	27.2%	0.87	0.42
Chrome_39	allreviews, priordefects	19.8%	0.75	0.5
Chrome_40	priordefects	25.4%	0.76	0.48
Chrome_41	allchangescount, priordefects	30.7%	0.78	0.51
Chrome_42	priordefects	32.9%	0.77	0.46
Chrome_43	minorauthors	25.6%	0.76	0.42
Chrome_44	allchangescount, minorauthors	36.1%	0.8	0.48
Chrome_combined	minorauthors, priordefects	19.5%	0.76	0.41
All	minorauthors, priordefects, size	32.7%	0.83	0.45

high, between 0.75 and 0.87, and the models had reasonable variance explained, between 19.5% and 51.4%.

4.3 Variables Directly Affecting Post Release Defects

Table 4 shows the variables that directly influence the *defects* node in the Markov Blanket. Of the 11 review measures one is present in Qt 50, *changesnodisc*, and another, *allreviews*, is present in Chrome release 39. In Table 5, we see the probability of a defect given each the review measure. A directory with two or more changes that are reviewed without discussion increases the probability of post-release defects by 42.3% in Qt 50. For Chrome 39, the more often a directory is reviewed the more often there are defects and in the case of 36 or more reviews for a directory the probability of observing a defect increases to 69.8%.

Of the 10 non-review “traditional” measures six are present in one or more releases, see Table 4. *allauthors*, *size* are present in one release each. *allchangescount*, and *minoraauthors*, are each in 2/8 releases. *priordefects* the most common predictor is present in 5/8 releases. In Table 5, we see the probability of a defect given each measure. We observe that all the predictors have a positive impact on defects, *i.e.* larger numbers increase the probability of defects. Below we discuss “traditional” measures that are present in more than one release. Having one or more *minoraauthors* that modify the files in a directory can also drastically increase the number of post-release defects by over 70%. Moderate changes to a component, *e.g.*, *allchangescount* < 50, can increase the number of defects by up to 33.8%. Components with many changes, > 50 can see post-release probabilities increase by over 60%. The relationship between *priordefects* is clearly shown in the CPT tables with a large number of prior defects, increasing the probability of post-release defects by between 30.6% and 73.7%.

4.4 Indirect Relationships and Visual Representation

In Figure 5, we present a snapshot of the BN model with only the *defects* node, the nodes in its Markov Blanket, and the nodes that directly affect the nodes in the Markov Blanket for ease of interpretation, since the complete BN models with all nodes are rather complicated. However, the complete BN models for individual Chrome and Qt releases, the aggregate Chrome and Qt datasets, and for the combined dataset are available in our replication package (Replication package 2018).

The dotted edges indicate that the coefficient is negative for that edge, *i.e.* increasing the value of the parent node decreases the value of the child node and vice versa. The immediate parents of the *defects* node (consequently, the Markov blanket for the *defects* node in this case) are colored in light blue and the *defects* node is denoted in a rectangular shape. The effect of each individual variable on the *defects* node is shown in Table 5. We do see a review measure: *allreviews* affecting the variable *minoraauthors* in the BN, but it has no *direct* impact on the number of defects in our combined model.

The variables indicated as the most important in the traditional modeling approach are also the ones indicated as most important in BN modeling approach, except in the traditional model the *minor authors* variable was discarded as being redundant, and *all authors* was used instead, while in BN approach *minor authors* have a direct influence on *defects*. This illustrates the ability of BNs to address some of the issues posed by correlated predictors, a situation common in software engineering.

The results from the two approaches are largely consistent in terms of indicating which variables are most significant in explaining the post release defects, and both approaches

Table 5 CPT: The conditional probability of post-release defects for each measure. As an example, having one or more minor authors making changes in a component increases the probability of post-release defects by over 70%

For release Qt 50: 12.5% of observations had defect

priordefects	probability of defects	changesnodisc	probability of defects
0	4.8%	0	5.6%
1	7.6%	1	9.7%
[2, 528]	30.6%	[2, 95]	42.3%

For release Qt 51: 17.1 % of observations had defect

allauthors	probability of defects	size	probability of defects
1	6.1%	[0, 78)	4.1%
2	10.5%	[78, 395)	11.7%
[3, 57]	44.1%	[395, 81654]	34.9%

For Qt - Combined: 14.4 % of observations had defect

priordefects	probability of defects	changesnodisc	probability of defects
0	7.0%	0	7.3%
1	10.4%	1	14.4%
[2, 624]	30.8%	[2, 191]	40.9%

For release Chrome 39: 42.1% of observations had defect

priordefects	probability of defects	allreviews	probability of defects
[0, 5)	20.7%	[0, 10)	18.3%
[5, 18)	45.1%	[10, 36)	39.4%
[18, 1588]	64.9%	[36, 1573]	69.8%

For release Chrome 40: 32.2% of observations had defect

priordefects	probability of defects
[0, 5)	11.9%
[5, 18)	28.4%
[18, 1588]	59.6%

For release Chrome 41: 39.8% of observations had defect

priordefects	probability of defects	allchangescount	probability of defects
[0, 4)	16.0%	[1, 10)	13.5%
[4, 18)	33.5%	[10, 48)	33.8%
[18, 1660]	73.7%	[48, 2415]	73.4%

For release Chrome 42: 38.8% of observations had defect

priordefects	probability of defects
[0, 4)	15.3%
[4, 17)	34.9%
[17, 1649]	69.2%

For release Chrome 43: 28.8% of observations had defect

minoraauthors	probability of defects
0	19.1%
[1, 108]	71.4%

Table 5 (continued)

For release Qt 50: 12.5% of observations had defect

priordefects	probability of defects	changesnodisc	probability of defects
--------------	------------------------	---------------	------------------------

For release Chrome 44: 30% of observations had defect

allchangescount	probability of defects	minorauthors	probability of defects
[1, 13)	8.1%	0	19.6%
[13, 52)	23.7%	[1, 96]	75.9%
[52, 1925]	59.4%		

For Chrome - Combined: 35.6% of observations had defect

priordefects	probability of defects	minorauthors	probability of defects
[0, 4)	13.8%	0	25.2%
[4, 18)	32.0%	[1, 108]	76.1%
[18, 1702]	63.1%		

For All releases combined: 26.9% of observations had defect

priordefects	probability of defects	minorauthors	probability of defects
[0, 2)	7.4%	0	18.9%
[2, 8)	23.1%	[1, 108]	69.5%
[8, 1702]	54.6%		
size	probability of defects		
[0, 113)	7.3%		
[113, 571)	21.9%		
[571, 106595]	50.3%		

For release Chrome 44: 30% of observations had defect

allchangescount	probability of defects	minorauthors	probability of defects
[1, 13)	8.1%	0	19.6%
[13, 52)	23.7%	[1, 96]	75.9%
[52, 1925]	59.4%		

For Chrome - Combined: 35.6% of observations had defect

priordefects	probability of defects	minorauthors	probability of defects
[0, 4)	13.8%	0	25.2%
[4, 18)	32.0%	[1, 108]	76.1%
[18, 1702]	63.1%		

For All releases combined: 26.9% of observations had defect

priordefects	probability of defects	minorauthors	probability of defects
[0, 2)	7.4%	0	18.9%
[2, 8)	23.1%	[1, 108]	69.5%
[8, 1702]	54.6%		
size	probability of defects		
[0, 113)	7.3%		
[113, 571)	21.9%		
[571, 106595]	50.3%		

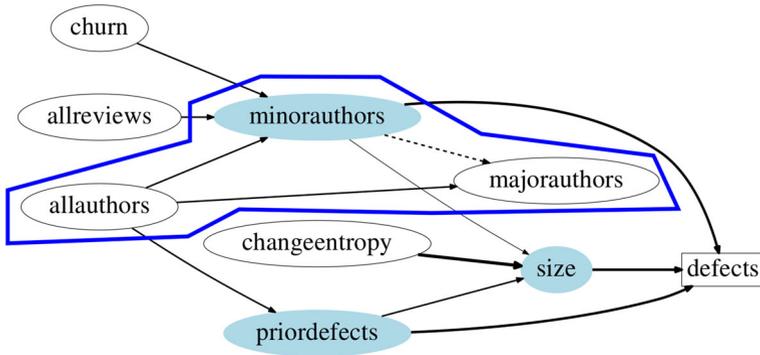


Fig. 5 The Bayesian network graph for the combined Chrome and Qt data. The model confirms that review measures only indirectly impact post-release defects. Traditional measures, such as prior defects, have direct impact

show that review-related measures have no direct influence over the post release defects variable in the combined model and no review measures are statistically significant in more than one individual release. By obtaining the same result using two completely independent modeling approaches increases our confidence in the findings. Our conclusion to RQ 3, Structure of the relationships, is summarized in the following box.

Conclusion 3: Only prior defects, module size, and minor authors have direct effects on (form a Markov blanket for) post-release defects in the combined model

4.5 Addressing the Issue of Highly Correlated Variables — Problem of Subjectivity in Variable Selection

We have claimed before that BN modeling approach is not affected by the presence of highly correlated variables, and that can be seen in our BN model as well. The three author related variables: *all authors*, *minor authors*, and *major authors* were highly correlated in our data. Therefore, in the traditional modeling approach only one of them, *all authors* was used in the final model.

In the combined BN model, the three variables appear connected to each other, as can be seen in Fig. 5 (the three nodes are inside the blue dotted polygon). The relationship among the nodes from the BN model is easily interpretable: more *allauthors* implies more *minor authors* and *major authors*, while increase in *minor authors* inevitably decreases *major authors* as *all authors* is the sum of minor and major authors. The BN model also suggests that the *minor authors* variable has substantially more influence over *defects* than *major authors* or *all authors*, thus it resolves the subjectivity in the variable selection problem.

To illustrate the usefulness of BNs it is worth making a few additional observations. The size of the module tends to be associated with code smells, effort, and defects (see, e.g., (Sjoberg et al. 2013; Nagappan et al. 2008; Arisholm and Briand 2006; Mockus 2010)). Not surprisingly, size affects both prior defects and defects, since relative module size tends to be stable release to release.

More interestingly, the BN model in Figure 5 suggests that, for example, the presence of *minor authors* both, increases the size of the module (perhaps via unnecessary code bloat), and also has a direct effect on the number of post-release defects (perhaps due to lack of understanding of the module). Thus it has a double effect on defects: direct, and mediated via module size.

The variable *minor authors* is, in turn, affected by the total number of authors, the number of changes made to the module, and the number of review issues. Arguably, the arrow should be pointing towards the review issues from minor authors as it is the minor authors that are likely to submit problematic code or be screened more vigorously during the review (see more discussion on incorporation prior knowledge in Section 5). However all these three relationships (except the direction of the third, which can be addressed by introducing a suitable prior) are rather intuitive.

Finally, it is worth considering the most important predictor of defects: prior defects. Apart from size, it is also related to the proportion of changes with no discussion, suggesting less aggressive reviews, the typical number of reviewers, and, surprisingly, is better for more complex modules. As noted earlier, the arrows should arguably be reversed: modules with prior defects probably invite more scrutiny with a larger review team. Why the modules with larger review teams tend to have higher proportion of changes with no discussion may be worth a further investigation.

Conclusion 4: *BN's can help address some difficulties posed by correlated predictors and the generative models help articulate potential mechanisms of how development process and product measure interact.*

5 Limitations

In this section, we discuss factors that in our opinion may pose a threat to validity of the results we present. We inherit validity threats from the study we replicate and discuss new threats related to BNs.

5.1 External Validity

McIntosh et al. (McIntosh et al. 2016) mined Android, LibreOffice, QT, ITK, and VTK, but in the end the only project with enough bugs and links to reviews for a reasonable analysis was QT. Given the instability of the predictors and the difficulty in linking reviews on projects, we decided to use a new modelling framework on a single large successful project, Google Chrome, instead of a broad study of the predictors across multiple projects. While single case studies have value (Runeson and Höst 2009b; Menzies et al. 2013), clearly our results do not generalize beyond these two projects.

5.2 Regression Model

We remove highly correlated variables and use the same model as McIntosh et al. (McIntosh et al. 2014). We do not consider interactions among variables because the model was already unstable and the additional complexity would further reduce stability in variable selection.

5.3 Latent Variables

Dealing with hidden variables in Bayesian Networks remains an open research question and an inherent limitation to all modeling techniques dealing with real observational data. However, this problem is not a serious threat to our results, since we do not attempt to establish any causal relationship among the variables. Our assumption to exclude potentially relevant unobserved variables is ameliorated by the use of prominent predictors of software defects used in extensive prior research on the subject.

5.4 Discretization

We transform our count variables to discrete variables using the Equal Frequency method as discussed in Section 4.1.1, and use two or three levels, based on the distribution of the original variables, for our discretized variables for the sake of simplicity in our final model. No discretization method is optimal, and the choice of the number of levels has subjectivity. However, as can be seen in the with the conditional probabilities in Table 5 the interpretations and bins seem reasonable given the software engineering context.

5.5 Threshold

In order to obtain the final structure from averaged model we use an arbitrary threshold of confidence. We verify the robustness of the network by gradually reducing the threshold and plotting the new structure. The conclusion of the sensitivity analysis is that the overall structure remains stable. In particular, the Markov Blanket of *defects* variable remains unchanged even for a threshold value of 0.45.

5.6 False Positive/Negative Edges in the Bayesian Network

We used the best performing BN structure method as reported in (Dey and Mockus 2020), and the final models were constructed based on repeating the search process 100 times. However, there is still the possibility of false positive or negative edges in the model, but the impact of it on the final result is unlikely to be significant.

5.7 Prior Knowledge in BN Structure Search

We do not use any prior knowledge of the problem domain while learning the BN structure. For instance, we have some prior knowledge about the directions: *e.g.*, the *defects* node should not have any outgoing edges since it is measured after the release, and the *prior defects* node should not have any incoming edges since this information is known a-priori. This knowledge can be incorporated into the search process by providing the initial partial structure as a parameter for the search function. Our unrestricted structure search yielded a model where the first assumption does hold, but second one does not. There is room for an argument that incorporating this prior knowledge will result in a more realistic model, but a counter-argument may be made as well. For example, in our model *prior defects* might represent a proxy measure for the inherent defectiveness of the module, and using the assumed prior knowledge would have excluded this possibility. Since this analysis was primarily concerned with direct effects on defects and all the discovered links were pointing inward (rendering the question moot), the ways to specify and incorporate expert knowledge while being important by itself is beyond the scope of this analysis.

6 Discussion of Related Work

In 1976 Fagan published the first empirical evaluation of software review, *i.e.* inspection (Fagan 1976). The work quantified the defect finding effectiveness of inspection based on the number of defects found per thousand lines of source code (KLOC) and percentage of total defects found by inspection. On the IBM system under study 38 defects per KLOC were found by inspection vs 8 per KLOC found by unit tests. Inspection found 82% of the total defects found for the released product. In the intervening 40 years, code review has changed dramatically from the rigid inspection process that Fagan introduced.

Most of the early work on inspection focused on minor variations in the inspection process but kept the formality, measurability, and rigidity intact (Martin and Tsai 1990; Knight and Myers 1993; Kollanus and Koskinen 2009; Laitenberger and DeBaud 2000; Wiegers 2001). The most important finding was that the inspection meeting need not be held in person to find a substantial number of defects (Votta 1993; Eick et al. 1992; Perry et al. 2002). This led the way to online review tools that ultimately lead to the currently popular and widely studied Gerrit (Mukadam et al. 2013; Mcintosh et al. 2016) and the pull request mechanism of GitHub (Yu et al. 2015; Rahman and Roy 2014; Gousios et al. 2015).

There is also a long history of examining the factors that make peer review effective. Porter et al. (Porter et al. 1998a) examined both the process and the inputs to the process (*e.g.*, reviewer expertise, and artifact complexity). In terms of the number of defects found during review, Porter et al. concluded that the best predictor was the level of expertise of the reviewers. Varying the processes had a negligible impact on the number of defects found. This finding is echoed by others (*e.g.*, (Sauer et al. 2000b; Kollanus and Koskinen 2009)).

Rigby et al. (Rigby et al. 2008; Rigby et al. 2012; Rigby and Storey 2011; Rigby et al. 2014) examined open source software based review on multiple projects including the Linux kernel, the Apache server, and KDE. They created regression models with the number of defects found during review and the amount of time take for review. They found remarkably similar practices across project that had very little process, but relied on expert reviewers frequently reviewing each commit. In a study at Microsoft and AMD, Rigby and Bird (Rigby and Bird 2013) found that these lightweight review practices were also used in industry. They also found that the focus had shifted from a defect finding activity to a problem solving one.

Recent works have focused on the non-defect finding benefits of code review. For example, interviews of Microsoft and OSS developers have been conducted to understand developer motivations for code review (Bacchelli and Expectations 2013; Bosu et al. 2017). They found that while developers want to find defects, they were also interested in spreading knowledge and discussing alternative solutions. Indeed, code review has also been shown to be effective at spreading knowledge and reducing the impact of code ownership (Rigby and Bird 2013; Kononenko et al. 2015; Thongtanunam et al. 2016). Other works focused on the types and utility of feedback provided by developers (Bosu et al. 2015; Beller et al. 2014; Kononenko et al. 2016) and on the ability of code reviews to identify security vulnerabilities (Bosu et al. 2014; Munaiah et al. 2017)

Despite these additional benefits of code review, the primary goal is still defect finding (Bacchelli and Expectations 2013; Bosu et al. 2017). The literature abounds with papers that use product and process metrics to predict where defects will occur, for example, (Fenton and Neil 1999a; Neuhaus et al. 2007; Shivaji et al. 2013). These models have also been used to understand changes in development practices, such as co-location vs remote developers (Bird et al. 2011), the impact of developer turnover (Mockus 2010) and much more. As far as we know, McIntosh *et al.*'s (Mcintosh et al. 2016; McIntosh et al. 2014) is the first

to examine to include peer review measures into a defect model. Earlier works (Porter et al. 1998a; Rigby et al. 2014) measured how many defects were found during the review, but did not look at the long-term impact of review on defects. As a result, our work first replicates McIntosh *et al.*'s work that covered only releases (two Qt releases, one release from ITK, and one from VTK), we expand the study to include six releases of the Chrome project.

A case for use of BNs in the context of Software Engineering was made by Fenton et al. (Fenton and Neil 1999b; Fenton et al. 2002), while the earliest publications utilizing BNs we could find (Herbsleb and Mockus 2003) constructed search of the structure based on the statistical significance of partial correlations in the context of modeling delays in globally distributed development. (Stamelos et al. 2003; Pendharkar et al. 2005) considered the application of Bayesian networks to prediction of effort, (Fenton et al. 2007; Neil and Fenton 1996; Okutan and Yildiz 2014) used Bayesian networks to predict defects, and (Pai and Dugan 2007) used BN approach for an empirical analysis of faultiness of a software. In a similar work, (Bai 2005) used modified BNs (Markov Bayesian network) for software reliability prediction. (Van Koten and Gray 2006) used BNs for predicting maintainability of Object Oriented software, and (Bibi et al. 2003) used BNs as a software productivity estimation tool. We are not aware of prior applications of Bayesian Networks for modeling software reviews. On the other hand, Bayesian structure learning is a big domain in itself with a wide range of algorithms, but its use in software engineering context is not very common.

6.1 Conclusion

Prior works have shown that the defects are both effectively and efficiently found during code review (Fagan 2002; Porter et al. 1998b; Rigby et al. 2014). Recent works provided qualitative evidence that reviews provide benefits beyond defect detection, such as knowledge sharing (Sauer et al. 2000a; Bacchelli and Expectations 2013; Rigby and Storey 2011; Bosu et al. 2015; Kononenko et al. 2016; Rahman et al. 2017). In contrast, the goal of this work is to understand if code review measures can quantify the longterm impact of peer review on post-release defects.

6.1.1 Conclusion 1: Reproduction and Replication

McIntosh et al. (McIntosh et al. 2014; McIntosh et al. 2016) were the first to study the impact of code review measures on post-release defects. We replicated their study using data they provided and as well as on the Chrome data we extracted. We discussed our findings with the first author of the original study. McIntosh et al. found that review participation had an influence on post-release defects, but we were unable to replicate these results. Instead we found that review measures contributed little to the performance of the model. The R^2 values with and without review measures were almost identical. In agreement with existing defect prediction work (Mockus et al. 2000; Hassan 2009; Bird et al. 2011; Graves et al. 2000), our results show that prior defects, the module size, and the number of authors are the strongest predictors of post-release defects. Review measures are neither necessary nor sufficient to create a good defect prediction model.

6.1.2 Conclusion 2: Inconsistent Models

It is extremely difficult to replicate an empirical software study that involves both mining operational data and statistical modelling. Despite using exactly the same data and

modelling approach we obtain substantially different results. In both our study and that of McIntosh et al. (McIntosh et al. 2016) a key problem is the need to select an uncorrelated set of variables. The variable selection process is inherently subjective because differences in expert opinions may lead to different sets of variables.

Furthermore, in both studies, the models were performed per project and per release. Even strong predictors, such as prior-defects varied substantially in their predictive power between project releases. This result suggests an issue with the traditional variable selection used in regression models.

6.1.3 Conclusion 3: Direct Effects

Regression models require the researcher to define a response and a set of predictors. This approach lacks tools to distinguish between an actual relationship and the effect of a shared confound. In contrast, Bayesian Networks remove the need for variable selection and shows the Bayesian relationships among variables. The term “direct effect” is meant to quantify an influence that is not mediated by other variables in the model or, more accurately, the sensitivity of Y to changes in X while all other factors in the analysis are held fixed. Indirect effects can manifest themselves on the response only through affecting the value of predictors that gave direct effects on the response.

According to our BN, only three measures directly impact post-release defects: the number of prior defects, the number of minor authors, and the size of the module. The code review measures, such as *rushed reviews*, *number of review participants*, and *discussion length*, did not directly impact the number of post-release defects.

6.1.4 Conclusion 4: Generative Models and Indirect Effects

The use of BN provides a way to evaluate the indirect effects that code reviews have on defects through the influence on other variables. Such indirect effects bedevil traditional analysis methods that use observational data. If the set of observed variables is complete, it is possible to calculate an impact of intervention akin to the results that could be obtained only in randomized experiments. For example, changes that have *no review discussion* tend to be associated with files that have had many *prior defects* which in turn increase the number of post-release defects. A further example from our BN model shows that having 5 or more reviewers is seen to increase chance of having post-release defects from 20% to 33% through mediating variables *allauthors* and *minorauthors*.

We have demonstrated the difficulties in using traditional models on observational data. Although individual code reviews find defects, we were unable to find any direct effect of review measures on post-release defects. By using BN we found that code review measures indirectly effect post-release defects. We hope that other researchers will use the approaches presented here to untangle the relationships among software measures. These indirect effects should provide a more nuanced understanding of software engineering. We make our scripts and data available in our replication package (Replication package 2018).

References

Almqvist JPF (2006) Replication of controlled experiments in empirical software engineering—a survey

- Arcuri A, Briand L (2011) A practical guide for using statistical tests to assess randomized algorithms in software engineering. In: 2011 33rd International Conference on Software Engineering (ICSE), pages 1–10. IEEE
- Arisholm E, Briand LC (2006) Predicting fault-prone components in a java legacy system. In: International Symposium on Empirical Software Engineering, pp 8–17
- Austin P, Tu J (2004) Automated variable selection methods for logistic regression result in unstable models for predicting ami mortality. *Journal of clinical epidemiology* 57:1138–46, 12
- Axelrod R (1997) Advancing the art of simulation in the social sciences. In: *Simulating social phenomena*, pages 21–40. Springer
- Bacchelli A, Expectations CB (2013) Outcomes, and challenges of modern code review. In: *Proceedings of the International Conference on Software Engineering*, pages 712–721, IEEE Press, 2013
- Bai CG (2005) Bayesian network based software reliability prediction with an operational profile. *J Syst Softw* 77(2):103–112
- Beller M, Bacchelli A, Zaidman A, Juergens E (2014) Modern code reviews in open-source projects: Which problems do they fix? In: *Proceedings of the 11th Working Conference on Mining Software Repositories, MSR*, pages 202–211, New York, NY, USA, 2014. ACM.
- Bibi S, Stamelos I, Angelis L (2003) Bayesian belief networks as a software productivity estimation tool. In: *1st Balkan Conference in Informatics*, Thessaloniki, Greece
- Bird C, Nagappan N, Murphy B, Gall H, Devanbu P (2011) Don't touch my code!: examining the effects of ownership on software quality. In: *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*, pages 4–14. ACM
- Bosu A, Carver JC, Bird C, Orbeck J, Chockley C (2017) Process aspects and social dynamics of contemporary code review: Insights from open source development and industrial practice at microsoft. *IEEE Trans Softw Eng* 43(1):56–75
- Bosu A, Carver JC, Hafiz M, Hilley P, Janni D (2014) Identifying the characteristics of vulnerable code changes: An empirical study. In: *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2014*, pages 257–268, New York, NY, USA, ACM
- Bosu A, Greiler M, Bird C (2015) Characteristics of useful code reviews: an empirical study at microsoft. In: *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*, pp 146–156
- Camilo F, Meneely A, Nagappan M (2015) Do bugs foreshadow vulnerabilities? a study of the chromium project. In: *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*, pp 269–279
- Carvalho AM (2009) Scoring functions for learning bayesian networks. *Inesc-id Tec, Rep*, pp 12
- Carver R (1978) The case against statistical significance testing. *Harv Educ Rev* 48(3):378–399
- Chlebus BS, Nguyen SH (1998) On finding optimal discretizations for two attributes. In: *International Conference on Rough Sets and Current Trends in Computing*, pages 537–544. Springer
- Dey T, Mockus A (2020) Deriving a usage-independent software quality metric. *Empir Softw Eng* 25(2):1596–1641
- Eick SG, Loader CR, Long MD, Votta LG, Wiel SV (1992) Estimating software fault content before coding. In: *Proceedings of the 14th International Conference on Software Engineering*, pp 59–65
- Fagan M (2002) A history of software inspections. In: *Software pioneers*, pages 562–573. Springer
- Fagan ME (1976) Design and code inspections to reduce errors in program development. *IBM Syst J* 15(3):182–211
- Fenton N, Krause P, Neil M (2002) Software measurement: Uncertainty and causal modeling. *IEEE software* 19(4):116–122
- Fenton N, Neil M, Marsh W, Hearty P, Marquez D, Krause P, Mishra R (2007) Predicting software defects in varying development lifecycles using bayesian nets. *Information and Software Technology* 49(1):32–43
- Fenton NE, Neil M (1999) A critique of software defect prediction models. *IEEE Trans Softw Eng* 25(5):675–689
- Fenton NE, Neil M (1999) A critique of software defect prediction models. *IEEE Transactions on software engineering* 25(5):675–689
- Fleiss JL, Levin B, Paik MC et al (1981) The measurement of interrater agreement. *Statistical methods for rates and proportions* 2(212-236):22–23
- Friedman N, Goldszmidt M, Wyner A (1999) Data analysis with bayesian networks: A bootstrap approach. In: *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 196–205. Morgan Kaufmann Publishers Inc.
- Garcia S, Luengo J, Sáez JA, Lopez V, Herrera F (2013) A survey of discretization techniques: Taxonomy and empirical analysis in supervised learning. *IEEE Transactions on Knowledge and Data Engineering* 25(4):734–750
- Gómez OS, Juristo N, Vegas S (2014) Understanding replication of experiments in software engineering: a classification. *Inf Softw Technol* 56(8):1033–1048

- Gousios G, Zaidman A, Storey M-A, van Deursen A (2015) Work practices and challenges in pull-based development: The integrator's perspective. In: Proceedings of the 37th International Conference on Software Engineering - Volume 1, ICSE '15, pages 358–368, Piscataway, NJ, USA, IEEE Press
- Graves TL, Karr AF, Marron JS, H. Siy (2000) Predicting fault incidence using software change history. *Software Engineering IEEE Transactions on* 26(7):653–661
- Harrell Jr. FE (2013) rms: Regression modeling strategies. r package version 4.0-0 City
- Hassan AE (2009) Predicting faults using the complexity of code changes. In: Proceedings of the 31st International Conference on Software Engineering, pages 78–88 IEEE Computer Society
- Heckerman D (1998) A tutorial on learning with bayesian networks. In: *Learning in graphical models*, pages 301–354. Springer
- Herbsleb JD, Mockus A (2003) An empirical study of speed and communication in globally-distributed software development. *IEEE Trans Softw Eng* 29(6):481–494
- Højsgaard S et al (2012) Graphical independence networks with the grain package for r. *J Stat Softw* 46(10):1–26
- Huang L, Boehm B (2006) How much software quality investment is enough: a value-based approach. *IEEE software* 23(5):88–95
- James G, Witten D, Hastie T, Tibshirani R (2013) *An introduction to statistical learning*, volume 112 Springer
- Knight JC, Myers EA (1993) An improved inspection technique. *ACM Communications* 36(11):51–61
- Kollanus S, Koskinen J (2009) Survey of software inspection research. *Open Software Engineering Journal* 3:15–34
- Koller D, Friedman N (2009) *Probabilistic graphical models: principles and techniques*
- Kononenko O, Baysal O, Godfrey MW (2016) Code review quality: How developers see it. In: 2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE), pp 1028–1038
- Kononenko O, Baysal O, Guerrouj L, Cao Y, Godfrey MW (2015) Investigating code review quality: Do people and participation matter? In: *Software Maintenance and Evolution (ICSME), 2015 IEEE International Conference on*, pages 111–120. IEEE
- Laitenberger O, DeBaud J (2000) An encompassing life cycle centric survey of software inspection. *J Syst Softw* 50(1):5–31
- Landis JR, Koch GG (1977) An application of hierarchical kappa-type statistics in the assessment of majority agreement among multiple observers. *Biometrics* 33:363–374
- Lauritzen SL, Spiegelhalter DJ (1988) Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society., Series B (Methodological)* 50(2):157–194
- Lin M, Lucas JrH. C., Shmueli G (2013) Research commentary—too big to fail: large samples and the p-value problem. *Inf Syst Res* 24(4):906–917
- Martin J, Tsai WT (1990) N-fold inspection: a requirements analysis technique. *ACM Communications* 33(2):225–232
- McIntosh S, Kamei Y, Adams B, Hassan AE (2014) The impact of code review coverage and code review participation on software quality: A case study of the qt, vtk, and itk projects. In: *Proceedings of the 11th Working Conference on Mining Software Repositories*, pages 192–201 ACM
- Mcintosh S, Kamei Y, Adams B, Hassan AE (2016) An empirical study of the impact of modern code review practices on software quality. *Empirical Softw. Engg.* 21(5):2146–2189
- Menzies T, Brady A, Keung J, Hihn J, Williams S, El-Rawas O, Green P, Boehm B (2013) Learning project management decisions: a case study with case-based reasoning versus data farming. *IEEE Trans Softw Eng* 39(12):1698–1713
- Mockus A (2010) Organizational volatility and its effects on software defects. In: *ACM SIGSOFT / FSE*, pages 117–126, Santa Fe New Mexico, November, pp 7–11
- Mockus A (2014) Engineering big data solutions. In: *ICSE'14 FOSE*, pp 85–99
- Mockus A, Fielding RT, Herbsleb J (2000) A case study of open source software development: the apache server. In: *Proceedings of the 22nd international conference on Software engineering*, pages 263–272. Acm
- Morales R, McIntosh S, Khomh F (2015) Do code review practices impact design quality? a case study of the qt, vtk, and itk projects. In: *Software Analysis, Evolution and Reengineering (SANER), 2015 IEEE 22nd International Conference on*, pages 171–180. IEEE
- Mukadam M, Bird C, Rigby PC (2013) Gerrit software code review data from android. In: *2013 10th Working Conference on Mining Software Repositories (MSR)*, pp 45–48
- Munaiah N, Camilo F, Wigham W, Meneely A, Nagappan M (2017) Do bugs foreshadow vulnerabilities? an in-depth study of the chromium project. *Empir Softw Eng* 22(3):1305–1347
- Nagappan N, Murphy B, Basili VR (2008) The influence of organizational structure on software quality: an empirical case study. In: *ICSE, 2008*, pp 521–530

- Neil M, Fenton N (1996) Predicting software quality using bayesian belief networks. In: Proceedings of the 21st Annual Software Engineering Workshop, pages 217–230 NASA Goddard Space Flight Centre
- Neuhaus S, Zimmermann T, Holler C, Zeller A (2007) Predicting vulnerable software components. In: Inproceedings of the 14th ACM conference on Computer and communications security, pages 529–540 ACM
- Okutan A, Yıldız OT (2014) Software defect prediction using bayesian networks. *Empir Softw Eng* 19(1):154–181
- Pai GJ, Dugan JB (2007) Empirical analysis of software fault content and fault proneness using bayesian methods. *IEEE Transactions on software Engineering* 33(10):675–686
- Pear J (2014) Probabilistic reasoning in intelligent systems: networks of plausible inference
- Pendharkar PC, Subramanian GH, Rodger JA (2005) A probabilistic model for predicting software development effort. *IEEE Transactions on software engineering* 31(7):615–624
- Perez A, Larranaga P, Inza I (2006) Supervised classification with conditional gaussian networks:., Increasing the structure complexity from naive bayes. *International Journal of Approximate Reasoning* 43(1):1–25
- Perry D, Porter A, Wade M, Votta L, Perpich J (2002) Reducing inspection interval in large-scale software development. *Software Engineering, IEEE Transactions on* 28(7):695–705
- Pinheiro J, Bates D, DebRoy S, Sarkar D (2011) R development core team. 2010. nlme: linear and nonlinear mixed effects models. r package version 3.1-97. R Foundation for Statistical Computing Vienna
- Porter A, Siy H, Mockus A, Votta L (1998) Understanding the sources of variation in software inspections. *ACM Transactions Software Engineering Methodology* 7(1):41–79
- Porter A, Siy H, Mockus A, Votta LG (1998) Understanding the sources of variation in software inspections *ACM Transactions on Software Engineering and Methodology*
- Rahman MM, Roy CK (2014) An insight into the pull requests of github. In: proceedings of the 11th Working Conference on Mining Software Repositories, MSR 2014, pages 364–367, New York NY. USA. ACM
- Rahman MM, Roy CK, Kula RG (2017) Predicting usefulness of code review comments using textual features and developer experience. In: 2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR), pp 215–226
- Replication package (2018) Our scripts and data are available: <https://github.com/CESEL/ReviewPostReleaseDefectsReplication>
- Report N (2002) The economic impacts of inadequate infrastructure for software testing
- Rigby P, Cleary B, Painchaud F, Storey M-A, German D (2012) Contemporary peer review in action: Lessons from open source development. *IEEE software* 29(6):56–61
- Rigby PC, Bird C (2013) Convergent contemporary software peer review practices. In: Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, pages 202–212 ACM
- Rigby PC, German DM, Cowen L, Storey MA (2014) Peer review on Open-Source software projects: parameters, statistical models, and theory. *ACM Transactions on Software Engineering and Methodology* 23(4):35:1–35:33
- Rigby PC, German DM, Storey M-A (2008) Open Source Software Peer Review Practices: A Case Study of the Apache Server. In: ICSE '08: Proceedings of the 30th International Conference on Software engineering, pages 541–550, New York, NY, USA, ACM
- Rigby PC, Storey MA (2011) Understanding broadcast based peer review on open source software projects. In: Inproceedings of the 33rd International Conference on Software Engineering, pages 541–550 ACM
- Runeson P, Höst M (2009) Guidelines for conducting and reporting case study research in software engineering. *Empirical software engineering* 14(2):131
- Runeson P, Höst M. (2009) Guidelines for conducting and reporting case study research in software engineering. *Empirical Softw. Engg.* 14(2):131–164
- Sauer C, Jeffery DR, Land L, Yetton P (2000) The effectiveness of software development technical reviews:., a behaviorally motivated program of research. *IEEE Transactions on Software Engineering* 26(1):1–14
- Sauer C, Jeffery DR, Land L, Yetton P (2000) The Effectiveness of Software Development Technical Reviews:., A Behaviorally Motivated Program of Research. *IEEE Transactions Software Engineering* 26(1):1–14
- Scutari M (2013) Learning bayesian networks in r, an example in systems biology. <http://www.bnlearn.com/about/slides/slides-useRconf13.pdf>
- Shivaji S, Whitehead EJ, Akella R, Kim S (2013) Reducing features to improve code change-based bug prediction. *IEEE Trans Softw Eng* 39(4):552–569
- Shmueli G (2010) To explain or to predict?. *Statistical science*, pp 289–310
- Shull F, Basili V, Carver J, Maldonado JC, Travassos GH, Mendonça M, Fabbri S (2002) Replicating software engineering experiments: addressing the tacit knowledge problem. In: *Empirical Software Engineering, 2002. Proceedings. 2002 International Symposium n*, pages 7–16. IEEE

- Shull FJ, Carver JC, Vegas S, Juristo N (2008) The role of replications in empirical software engineering. *Empirical software engineering* 13(2):211–218
- Sjoberg DI, Yamashita A, Anda B, Mockus A, Dyba T (2013) Quantifying the effect of code smells on maintenance effort. *IEEE Trans Softw Eng* 39(8):1144–1156
- Sober E (2002) Instrumentalism, parsimony, and the akaike framework. *Philos Sci* 69(S3):S112–S123
- Stamelos I, Angelis L, Dimou P, Sakellaris E (2003) On the use of bayesian belief networks for the prediction of software productivity. *Inf Softw Technol* 45(1):51–60
- Thongtanunam P, McIntosh S, Hassan AE, Iida H (2016) Revisiting code ownership and its relationship with software quality in the scope of modern code review. In: *Proceedings of the 38th International Conference on Software Engineering, ICSE '16*, pages 1039–1050, New York, NY, USA, ACM
- Van Koten C, Gray A (2006) An application of bayesian network for predicting object-oriented software maintainability. *Inf Softw Technol* 48(1):59–67
- Votta LG (1993) Does every inspection need a meeting? *SIGSOFT Softw Eng. Notes* 18(5):107–114
- Wiegiers KE (2001) *peer reviews in software: a practical guide*. Addison-wesley information technology series Addison-Wesley
- Yu Y, Wang H, Filkov V, Devanbu P, Vasilescu B (2015) Wait for it: Determinants of pull request evaluation latency on github. In: *2015 IEEE/ACM 12th Working Conference on Mining Software Repositories*, pp 367–371
- ZCarver JC (2010) Towards reporting guidelines for experimental replications: A proposal. In: *1st international workshop on replication in empirical software engineering*, pages 2–5. Citeseer
- Zheng Q, Mockus A, Zhou M (2015) A method to identify and correct problematic software activity data: Exploiting capacity constraints and data redundancies. In: *ESEC/FSE'15*, pages 637–648, Bergamo, Italy, ACM

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Andrey Krutauz¹ · Tapajit Dey² · Peter C. Rigby¹  · Audris Mockus²

Andrey Krutauz
andrey.krutauz@encs.concordia.ca

Tapajit Dey
tdey2@vols.utk.edu

Audris Mockus
audris@utk.edu

¹ Concordia University Montreal, Montreal QC, Canada

² University of Tennessee Knoxville, Knoxville, Tennessee, USA